
Provable Non-convex Phase Retrieval with Outliers: Median Truncated Wirtinger Flow

Huishuai Zhang

Department of EECS, Syracuse University, Syracuse, NY 13244 USA

HZHAN23@SYR.EDU

Yuejie Chi

Department of ECE, The Ohio State University, Columbus, OH 43210 USA

CHI.97@OSU.EDU

Yingbin Liang

Department of EECS, Syracuse University, Syracuse, NY 13244 USA

YLIANG06@SYR.EDU

Abstract

Solving systems of quadratic equations is a central problem in machine learning and signal processing. One important example is phase retrieval, which aims to recover a signal from only magnitudes of its linear measurements. This paper focuses on the situation when the measurements are corrupted by arbitrary outliers, for which the recently developed non-convex gradient descent Wirtinger flow (WF) and truncated Wirtinger flow (TWF) algorithms likely fail. We develop a novel median-TWF algorithm that exploits robustness of sample median to resist arbitrary outliers in the initialization and the gradient update in each iteration. We show that such a non-convex algorithm provably recovers the signal from a near-optimal number of measurements composed of i.i.d. Gaussian entries, up to a logarithmic factor, even when a constant portion of the measurements are corrupted by arbitrary outliers. We further show that median-TWF is also robust when measurements are corrupted by both arbitrary outliers and bounded noise. Our analysis of performance guarantee is accomplished by development of non-trivial concentration measures of median-related quantities, which may be of independent interest. We further provide numerical experiments to demonstrate the effectiveness of the approach.

1. Introduction

Phase retrieval is a classical problem in machine learning, signal processing and optical imaging, where one aims to

recover a signal $\mathbf{x} \in \mathbb{R}^n$ from only observing the magnitudes of its linear measurements:

$$y_i = |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2, \quad i = 1, \dots, m.$$

It has many important applications such as X-ray crystallography (Drenth, 2007), but is known to be notoriously difficult due to the quadratic form of the measurements. Classical methods based on alternating minimization between the signal of interest and the phase information (Fienup, 1982), though computationally simple, are often trapped at local minima and lack rigorous performance guarantees.

Using the lifting trick, the phase retrieval problem can be reformulated as estimating a rank-one positive semidefinite matrix $\mathbf{X} = \mathbf{x}\mathbf{x}^T$ from linear measurements (Balan et al., 2006), to which convex relaxations into semidefinite programming are considered (Waldspurger et al., 2015; Candès et al., 2013; Chen et al., 2015; Demanet & Hand, 2014; Candès & Li, 2014; Li & Voroninski, 2013). In particular, when the measurement vectors \mathbf{a}_i 's are composed of i.i.d. Gaussian entries, Phaselift (Candès et al., 2013) perfectly recovers all $\mathbf{x} \in \mathbb{R}^n$ with high probability as long as the number m of measurements is on the order of n .

However, the computational cost of Phaselift becomes prohibitive when the signal dimension is large. Appealingly, a so-called Wirtinger flow (WF) algorithm based on gradient descent was recently proposed in (Candès et al., 2015; Soltanolkotabi, 2014) and shown to work remarkably well: it converges to the global optima when properly initialized using the spectral method. The truncated Wirtinger flow (TWF) algorithm (Chen & Candès, 2015) further improves WF by eliminating samples whose contributions to both the initialization and the search direction are excessively deviated from the *sample mean*, so that the behavior of each gradient update is well controlled. TWF is shown to converge globally at a geometric rate as long as m is on the order of n for i.i.d. Gaussian measurement vectors using a constant step size. Both WF and TWF algorithms have

been shown to be robust to bounded noise in the measurements.

However, the performance of WF and TWF can be very sensitive to outliers that take arbitrary values and can introduce anomalous search directions. Even for TWF, since the sample mean can be arbitrarily perturbed, the truncation rule based on such sample mean cannot control the gradient well. On the other hand, the ability to handle outliers is of great importance for phase retrieval algorithms, because outliers arise frequently from the phase imaging applications (Weller et al., 2015) due to various reasons such as detector failures, recording errors, and missing data. While a form of Phaselift (Hand, 2015) is shown to be robust to sparse outliers even when they constitute a constant portion of all measurements, it is computationally too expensive.

1.1. Main Contributions

The main contribution of this paper lies in the development of a non-convex phase retrieval algorithm with both statistical and computational efficiency, and provable robustness to even a constant proportion of outliers. To the best of the authors' knowledge, our work is the first application of the median to robustify high-dimensional statistical estimation in the presence of arbitrary outliers with rigorous non-asymptotic performance guarantees.

Our strategy is to carefully robustify the TWF algorithm by replacing the sample mean used in the truncation rule by its robust counterpart, the *sample median*. We refer to the new algorithm as median truncated Wirtinger flow (median-TWF). Appealingly, median-TWF does not require any knowledge of the outliers. The robustness property of median lies in the fact that the median cannot be arbitrarily perturbed unless the outliers dominate the inliers (Huber, 2011). This is in sharp contrast to the mean, which can be made arbitrarily large even by a single outlier. Thus, using the sample median in the truncation rule can effectively remove the impact of outliers and indeed, the performance of median-TWF can be provably guaranteed.

Statistically, the sample complexity of median-TWF is near-optimal up to a logarithmic factor when the measurement vectors are composed of i.i.d. Gaussian entries. We demonstrate that as soon as the number m of measurements is on the order of $n \log n$, median-TWF converges to the global optima, i.e. recovers the ground truth up to a global sign difference, even when the number of outliers scales *linearly* with m . Computationally, median-TWF converges at a geometric rate, requiring a computational cost of $O(mn \log 1/\epsilon)$ to reach ϵ -accuracy, which is linear in the problem size. Reassuringly, under the same sample complexity, median-TWF still recovers the ground truth when outliers are absent. It can therefore handle outliers in an *oblivious* fashion. Finally, median-TWF is also stable when the measurements are further corrupted by dense bounded noise besides outliers.

Our proof proceeds by first showing the initialization of median-TWF is close enough to the ground truth, and that the neighborhood of the ground truth, where the initialization lands in, satisfies certain *Regularity Condition* (Candès et al., 2015; Chen & Candès, 2015) that guarantees convergence of the descent rule, as long as the size of the corruption is small enough and the sample size is large enough. However, as a nonlinear operator, the sample median used in median-TWF is much more difficult to analyze than the sample mean used in TWF, which is a linear operator and many existing concentration inequalities are readily applicable. Considerable technical efforts lie in developing novel non-asymptotic concentrations of the sample median, and various statistical properties of the sample median related quantities, which may be of independent interest.

1.2. Related Work

The adoption of median in machine learning and computer science is not unfamiliar, for example, K -median clustering (Chen, 2006) and resilient data aggregation for sensor networks (Wagner, 2004). Our work here further extends the applications of median to robustifying high-dimensional estimation problems.

Another popular approach in robust estimation is to use the trimmed mean (Huber, 2011), which has found success in robustifying sparse regression (Chen et al., 2013), linear regression (Bhatia et al., 2015), subspace clustering (Qu & Xu, 2015), etc. However, using the trimmed mean requires knowledge of an upper bound on the number of outliers, whereas median does not require such information.

Developing non-convex algorithms with provable global convergence guarantees has attracted intensive research interest recently. A partial list includes low-rank matrix recovery (Keshavan et al., 2010; Zheng & Lafferty, 2015; Tu et al., 2015; Chen & Wainwright, 2015; White et al., 2015), robust principal component analysis (Netrapalli et al., 2014), robust tensor decomposition (Anandkumar et al., 2015), dictionary learning (Arora et al., 2015; Sun et al., 2015), etc. We expect our analysis in this paper may be extended to robustly solving other systems of quadratic or bilinear equations in a non-convex fashion, such as mixed linear regression (Chen et al., 2014), sparse phase retrieval (Cai et al., 2015), and blind deconvolution (Lee et al., 2015).

1.3. Paper Organization and Notations

The rest of this paper is organized as follows. Section 2 provides the problem formulation, and Section 3 describes the proposed median-TWF algorithm. Theoretical performance guarantees are stated in Section 4, with proof outlines given in Section 5. Numerical experiments are presented in Section 6. Finally, we conclude in Section 7.

We adopt the following notations in this paper. Given a vec-

tor of numbers $\{\beta_i\}_{i=1}^m$, the sample median is denoted as $\text{med}(\{\beta_i\}_{i=1}^m)$. The indicator function $\mathbb{1}_A = 1$ if the event A holds, and $\mathbb{1}_A = 0$ otherwise. For two matrices, $\mathbf{A} \prec \mathbf{B}$ if $\mathbf{B} - \mathbf{A}$ is a positive semidefinite matrix. We define the Euclidean distance between two vectors up to a global sign difference as $\text{dist}(\mathbf{z}, \mathbf{x}) := \min \{\|\mathbf{z} - \mathbf{x}\|, \|\mathbf{z} + \mathbf{x}\|\}$.

2. Problem Formulation

Suppose the following set of m measurements are given

$$y_i = |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2 + \eta_i, \quad i = 1, \dots, m, \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^n$ is the unknown signal,¹ $\mathbf{a}_i \in \mathbb{R}^n$ for $i = 1, \dots, m$ are measurement vectors with each \mathbf{a}_i having *i.i.d.* Gaussian entries distributed as $\mathcal{N}(0, 1)$, and $\eta_i \in \mathbb{R}$ for $i = 1, \dots, m$ are outliers with arbitrary values. We assume that outliers are sparse with sm nonzero values, i.e., $\|\boldsymbol{\eta}\|_0 \leq sm$, where $\boldsymbol{\eta} = \{\eta_i\}_{i=1}^m \in \mathbb{R}^m$. Here, s is a nonzero constant, representing the fraction of measurements that are corrupted by outliers.

We are also interested in the model when the measurements are corrupted by not only sparse arbitrary outliers but also dense bounded noise. Under such a model, the measurements are given by

$$y_i = |\langle \mathbf{a}_i, \mathbf{x} \rangle|^2 + w_i + \eta_i, \quad i = 1, \dots, m, \quad (2)$$

where the bounded noise $\mathbf{w} = \{w_i\}_{i=1}^m$ satisfies $\|\mathbf{w}\|_\infty \leq c_1 \|\mathbf{x}\|^2$ for some universal constant c_1 , and as before, the outlier satisfies $\|\boldsymbol{\eta}\|_0 \leq sm$.

The goal is to recover the signal \mathbf{x} (up to a global sign difference) from the measurements $\mathbf{y} = \{y_i\}_{i=1}^m$ and measurement vectors $\{\mathbf{a}_i\}_{i=1}^m$.

3. Median-TWF Algorithm

A natural idea is to recover the signal as a solution to the following optimization problem

$$\min_{\mathbf{z}} \sum_{i=1}^m -\ell(\mathbf{z}; y_i) \quad (3)$$

where $\ell(\mathbf{z}, y_i)$ is a likelihood function, e.g., using Gaussian or Poisson likelihood. Since the measurements are quadratic in \mathbf{x} , the objective function is non-convex. A typical gradient descent procedure to solve (3) proceeds as

$$\mathbf{z}^{(t+1)} = \mathbf{z}^{(t)} + \frac{\mu_t}{m} \sum_{i=1}^m \nabla \ell(\mathbf{z}^{(t)}; y_i), \quad (4)$$

where $\mathbf{z}^{(t)}$ denotes the t th iterate of the algorithm, and μ_t is the step size. By a careful initialization using the spectral

¹We focus on real signals here, but our analysis can be extended to complex signals.

method, the WF algorithm (Candès et al., 2015) using the gradient descent update (4) is shown to converge globally under the Gaussian likelihood (i.e., quadratic loss) function, as long as the number of measurements is on the order of $n \log n$ for *i.i.d.* Gaussian measurement vectors.

The TWF algorithm (Chen & Candès, 2015) improves WF in both initialization and the descent rule: only a subset of samples T_0 contributes to the spectral method, and only a subset of data-dependent and iteration-varying samples T_{t+1} contributes to the search directions:

$$\mathbf{z}^{(t+1)} = \mathbf{z}^{(t)} + \frac{\mu_t}{m} \sum_{i \in T_{t+1}} \nabla \ell(\mathbf{z}^{(t)}; y_i). \quad (5)$$

The idea is that through pruning, i.e., samples with gradient components $\nabla \ell(\mathbf{z}^{(t)}; y_i)$ being much larger than the sample mean are truncated so that each update is well controlled. This modification yields both statistical and computational benefits – under the Poisson loss function, TWF converges globally geometrically to the true signal with a constant step size with measurements at the order of n and with *i.i.d.* Gaussian measurement vectors.

However, if some measurements are corrupted by arbitrary-valued outliers as in (1), both WF and TWF can fail. This is because the gradient of the loss function typically contains the term $|y_i - |\mathbf{a}_i^T \mathbf{z}|^2|$. With y_i being corrupted by arbitrarily large η_i , the gradient can deviate the search direction from the signal arbitrarily. In TWF, since the truncation rule is based on the sample mean of the gradient, which can be affected significantly even by a single outlier, we cannot expect it to converge globally, particularly when the fraction of corrupted measurements is linear with the total number m of measurements, which is the regime we are interested in.

To handle outliers, our central idea is to prune the samples in both the initialization and each iteration via the *sample median* related quantities. Compared to the sample mean used in TWF, the sample median is much less affected even in the presence of a certain linear fraction of outliers, and is thus more robust to outliers with arbitrary values.

In the following, we describe our median-TWF in more details. We adopt the following Poisson likelihood function,

$$\ell(\mathbf{z}; y_i) = y_i \log |\mathbf{a}_i^T \mathbf{z}|^2 - |\mathbf{a}_i^T \mathbf{z}|^2, \quad (6)$$

which is motivated by the maximum likelihood estimation of the signal when the measurements are corrupted by Poisson distributed noise. We note that our analysis is also applicable to the quadratic loss function, but in order to compare more directly to TWF in (Chen & Candès, 2015), we adopt the Poisson likelihood function in (6).

Our median-TWF algorithm (summarized in Algorithm 1) contains the following main steps:

1. **Initialization:** We initialize $\mathbf{z}^{(0)}$ by the spectral method with a truncated set of samples, where the threshold is de-

Algorithm 1 Median Truncated Wirtinger Flow (Median-TWF)

Input: $\mathbf{y} = \{y_i\}_{i=1}^m, \{\mathbf{a}_i\}_{i=1}^m$;

Parameters: thresholds $\alpha_y, \alpha_h, \alpha_l$, and α_u , stepsize μ_t ;

Initialization: Let $\mathbf{z}^{(0)} = \lambda_0 \tilde{\mathbf{z}}$, where $\lambda_0 = \sqrt{\text{med}(\mathbf{y})/0.455}$ and $\tilde{\mathbf{z}}$ is the leading eigenvector of

$$\mathbf{Y} := \frac{1}{m} \sum_{i=1}^m y_i \mathbf{a}_i \mathbf{a}_i^T \mathbb{1}_{\{|y_i| \leq \alpha_y^2 \lambda_0^2\}}. \quad (7)$$

Gradient loop: for $t = 0 : T - 1$ do

$$\mathbf{z}^{(t+1)} = \mathbf{z}^{(t)} + \frac{2\mu_t}{m} \sum_{i=1}^m \frac{y_i - |\mathbf{a}_i^T \mathbf{z}^{(t)}|^2}{\mathbf{a}_i^T \mathbf{z}^{(t)}} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i}, \quad (8)$$

where

$$\begin{aligned} \mathcal{E}_1^i &:= \left\{ \alpha_l \|\mathbf{z}^{(t)}\| \leq |\mathbf{a}_i^T \mathbf{z}^{(t)}| \leq \alpha_u \|\mathbf{z}^{(t)}\| \right\}, \\ \mathcal{E}_2^i &:= \left\{ |y_i - |\mathbf{a}_i^T \mathbf{z}^{(t)}|^2| \leq \alpha_h K_t \frac{|\mathbf{a}_i^T \mathbf{z}^{(t)}|}{\|\mathbf{z}^{(t)}\|} \right\}, \\ K_t &:= \text{med} \left(\left\{ |y_i - |\mathbf{a}_i^T \mathbf{z}^{(t)}|^2| \right\}_{i=1}^m \right). \end{aligned}$$

Output \mathbf{z}_T .

terminated by the median of $\{y_i\}_{i=1}^m$. In comparison, WF does not truncate samples, and the truncation in TWF is based on the mean of $\{y_i\}_{i=1}^m$, which is not robust to outliers. As will be shown, as long as the portion of outliers is not too large, our initialization (7) is guaranteed to be within a small neighborhood of the ground truth signal.

2. Gradient loop: for each iteration $0 \leq t \leq T - 1$, comparing (4) and (8), median-TWF uses an iteration-varying truncated gradient given as

$$\nabla \ell_{tr}(\mathbf{z}^{(t)}) = 2 \sum_{i=1}^m \frac{y_i - |\mathbf{a}_i^T \mathbf{z}^{(t)}|^2}{\mathbf{a}_i^T \mathbf{z}^{(t)}} \mathbf{a}_i \mathbb{1}_{\mathcal{E}_1^i \cap \mathcal{E}_2^i}. \quad (9)$$

It is clear from the definition of the set \mathcal{E}_2^i (see Algorithm 1), that samples are truncated by the sample median of gradient components evaluated at the current iteration, as opposed to the sample mean in TWF.

We set the step size in the median-TWF to be a fixed small constant, i.e., $\mu_t = 0.2$. The rest of the parameters $\{\alpha_y, \alpha_h, \alpha_l, \alpha_u\}$ are set to satisfy

$$\begin{aligned} \zeta_1 &:= \max \left\{ \mathbb{E} \left[\xi^2 \mathbb{1}_{\{|\xi| < \sqrt{1.01}\alpha_l \text{ or } |\xi| > \sqrt{0.99}\alpha_u\}} \right], \right. \\ &\quad \left. \mathbb{E} \left[\mathbb{1}_{\{|\xi| < \sqrt{1.01}\alpha_l \text{ or } |\xi| > \sqrt{0.99}\alpha_u\}} \right] \right\}, \\ \zeta_2 &:= \mathbb{E} \left[\xi^2 \mathbb{1}_{\{|\xi| > 0.248\alpha_h\}} \right], \\ 2(\zeta_1 + \zeta_2) + \sqrt{8/\pi}\alpha_h^{-1} &< 1.99 \\ \alpha_y &\geq 3, \end{aligned} \quad (10)$$

where $\xi \sim \mathcal{N}(0, 1)$. For example, we set $\alpha_l = 0.3, \alpha_u = 5, \alpha_y = 3$ and $\alpha_h = 12$, and consequently $\zeta_1 \approx 0.24$ and $\zeta_2 \approx 0.032$.

4. Performance Guarantees of Median-TWF

In this section, we characterize the performance guarantees of median-TWF.

Theorem 1 (Exact recovery with sparse arbitrary outliers). *Consider the phase retrieval problem with sparse outliers given in (1). There exist constants $\mu_0, s_0 > 0, 0 < \rho, \nu < 1$ and $c_0, c_1, c_2 > 0$ such that if $m \geq c_0 n \log n, s < s_0, \mu \leq \mu_0$, then with probability at least $1 - c_1 \exp(-c_2 m)$, the median-TWF yields*

$$\text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \leq \nu(1 - \rho)^t \|\mathbf{x}\|, \quad \forall t \in \mathbb{N} \quad (11)$$

simultaneously for all $\mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$.

Theorem 1 indicates that median-TWF admits exact recovery for all signals in the presence of sparse outliers with arbitrary magnitudes even when the number of outliers scales linearly with the number of measurements, as long as the number of samples satisfies $m \gtrsim n \log n$. This is near-optimal up to a logarithmic factor.

Moreover, median-TWF converges at a geometric rate using a constant step size, with per-iteration cost $O(mn)$ (note that the median can be computed in linear time (Tibshirani, 2008)). To reach ϵ -accuracy, i.e., $\text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \leq \epsilon$, only $O(\log 1/\epsilon)$ iterations are needed, and the total computational cost is $O(mn \log 1/\epsilon)$, which is highly efficient.

Not surprisingly, Theorem 1 implies that median-TWF also performs well for the noise-free model, as a special case of the model with outliers. This justifies utilization of median-TWF in an oblivious situation without knowing whether the underlying measurements are corrupted by outliers.

Corollary 1 (Exact recovery for noise-free model). *Suppose that the measurements are noise-free, i.e., $\eta_i = 0$ for $i = 1, \dots, m$ in the model (1). There exist constants $\mu_0 > 0, 0 < \rho, \nu < 1$ and $c_0, c_1, c_2 > 0$ such that if $m \geq c_0 n \log n$ and $\mu \leq \mu_0$, then with probability at least $1 - c_1 \exp(-c_2 m)$, the median-TWF yields*

$$\text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \leq \nu(1 - \rho)^t \|\mathbf{x}\|, \quad \forall t \in \mathbb{N} \quad (12)$$

simultaneously for all $\mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$.

We next consider the model when the measurements are corrupted by both sparse arbitrary outliers and dense bounded noise. Our following theorem characterizes that median-TWF is robust to coexistence of the two types of noises.

Theorem 2 (Stability to sparse arbitrary outliers and dense bounded noises). *Consider the phase retrieval problem given in (2) in which measurements are corrupted*

by both sparse arbitrary and dense bounded noises. There exist constants $\mu_0, s_0 > 0$, $0 < \rho < 1$ and $c_0, c_1, c_2 > 0$ such that if $m \geq c_0 n \log n$, $s < s_0$, $\mu \leq \mu_0$, then with probability at least $1 - c_1 \exp(-c_2 m)$, median-TWF yields

$$\text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \lesssim \frac{\|\mathbf{w}\|_\infty}{\|\mathbf{x}\|} + (1 - \rho)^t \|\mathbf{x}\|, \quad \forall t \in \mathbb{N} \quad (13)$$

simultaneously for all $\mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$.

Theorem 2 immediately implies the stability of median-TWF for the model corrupted only by dense bounded noise.

Corollary 2. Consider the phase retrieval problem in which measurements are corrupted only by dense bounded noises, i.e., $\eta_i = 0$ for $i = 1, \dots, m$ in the model (2). There exist constants $\mu_0 > 0$, $0 < \rho < 1$ and $c_0, c_1, c_2 > 0$ such that if $m \geq c_0 n \log n$, $\mu \leq \mu_0$, then with probability at least $1 - c_1 \exp(-c_2 m)$, median-TWF yields

$$\text{dist}(\mathbf{z}^{(t)}, \mathbf{x}) \lesssim \frac{\|\mathbf{w}\|_\infty}{\|\mathbf{x}\|} + (1 - \rho)^t \|\mathbf{x}\|, \quad \forall t \in \mathbb{N} \quad (14)$$

simultaneously for all $\mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$.

Thus, Theorem 2 and Corollary 2 imply that median-TWF for the model with both sparse arbitrary outliers and dense bounded noises achieves the same convergence rate and the same level of estimation error as the model with only bounded noise. In fact, together with Theorem 1 and Corollary 1, it can be seen that applying median-TWF does not require the knowledge of the noise corruption models. When there indeed exist outliers, median-TWF achieves the same performance as if the outliers do not exist.

5. Proof Outlines of Main Results

In this section, we first develop a few statistical properties of median that will be useful for our analysis of performance guarantees, and then provide a proof sketch of our main results. The details of the proofs can be found in Supplemental Materials.

5.1. Properties of Median

We start by the definitions of the quantile of a population distribution and its sample version.

Definition 1 (Generalized quantile function). Let $0 < p < 1$. For a cumulative distribution function (CDF) F , the generalized quantile function is defined as

$$F^{-1}(p) = \inf\{x \in \mathbb{R} : F(x) \geq p\}. \quad (15)$$

For simplicity, denote $\theta_p(F) = F^{-1}(p)$ as the p -quantile of F . Moreover for a sample sequence $\{X_i\}_{i=1}^m$, the sample p -quantile $\theta_p(\{X_i\})$ means $\theta_p(\hat{F})$, where \hat{F} is the empirical distribution of the samples $\{X_i\}_{i=1}^m$.

Remark 1. We take the median $\text{med}(\{X_i\}) = \theta_{1/2}(\{X_i\})$ and use both notations interchangeably.

Next, we show that as long as the sample size is large enough, the sample quantile concentrates around the population quantile (motivated from (Charikar et al., 2002)), as in Lemma 1.

Lemma 1. Suppose $F(\cdot)$ is cumulative distribution function (i.e., non-decreasing and right-continuous) with continuous density function $F'(\cdot)$. Assume the samples $\{X_i\}_{i=1}^m$ are i.i.d. drawn from F . Let $0 < p < 1$. If $l < F'(\theta) < L$ for all θ in $\{\theta : |\theta - \theta_p| \leq \epsilon\}$, then

$$|\theta_p(\{X_i\}_{i=1}^m) - \theta_p(F)| < \epsilon \quad (16)$$

holds with probability at least $1 - 2 \exp(-2m\epsilon^2 l^2)$.

Lemma 2 bounds the distance between the median of two sequences.

Lemma 2. Given a vector $\mathbf{X} = (X_1, X_2, \dots, X_n)$, reorder them in a non-decreasing manner

$$X_{(1)} \leq X_{(2)} \leq \dots \leq X_{(n-1)} \leq X_{(n)}.$$

Given another vector $\mathbf{Y} = (Y_1, Y_2, \dots, Y_n)$, then one has

$$|X_{(k)} - Y_{(k)}| \leq \|\mathbf{X} - \mathbf{Y}\|_\infty, \quad (17)$$

for all $k = 1, \dots, n$.

Lemma 3 suggests that in the presence of outliers, one can lower and upper bound the sample median by neighboring quantiles of the corresponding clean samples.

Lemma 3. Consider clean samples $\{\tilde{X}_i\}_{i=1}^m$. If a fraction s of them are corrupted by outliers, one obtains contaminated samples $\{X_i\}_{i=1}^m$ which contain sm corrupted samples and $(1 - s)m$ clean samples. Then

$$\theta_{\frac{1}{2}-s}(\{X_i\}) \leq \theta_{\frac{1}{2}}(\{X_i\}) \leq \theta_{\frac{1}{2}+s}(\{X_i\}).$$

Finally, Lemma 4 is related to bound the median and density at the median for the product of two possibly correlated standard Gaussian random variables.

Lemma 4. Let $u, v \sim \mathcal{N}(0, 1)$ which can be correlated with the correlation coefficient $|\rho| \leq 1$. Let $r = |uv|$, and $\psi_\rho(x)$ represent the density of r . Denote $\theta_{\frac{1}{2}}(\psi_\rho)$ as the median of r , and the value of $\psi_\rho(x)$ at the median as $\psi_\rho(\theta_{1/2})$. Then for all ρ ,

$$0.348 < \theta_{1/2}(\psi_\rho) < 0.455,$$

$$0.47 < \psi_\rho(\theta_{1/2}) < 0.76.$$

5.2. Robust Initialization with Outliers

We show that the initialization provided by the median-truncated spectral method in (7) is close enough to the

ground truth, i.e. $\text{dist}(\mathbf{z}^{(0)}, \mathbf{x}) \leq 1/11\|\mathbf{x}\|$, even if there are sm arbitrary outliers, as long as s is a small enough constant.

1. We first bound the concentration of $\text{med}(\{y_i\})$, also denoted by $\theta_{1/2}(\{y_i\})$. Lemma 3 suggests that

$$\theta_{\frac{1}{2}-s}(\{(\mathbf{a}_i^T \mathbf{x})^2\}) < \theta_{1/2}(\{y_i\}) < \theta_{\frac{1}{2}+s}(\{(\mathbf{a}_i^T \mathbf{x})^2\})$$

Observe that $(\mathbf{a}_i^T \mathbf{x})^2 = \tilde{a}_{i1}^2 \|\mathbf{x}\|^2$, where $\tilde{a}_{i1} = \mathbf{a}_i^T \mathbf{x} / \|\mathbf{x}\|$ is a standard Gaussian random variable. Thus $|\tilde{a}_{i1}|^2$ is a χ_1^2 random variable, whose CDF is denoted as K . By Lemma 1, for a small ϵ , one has $\left| \theta_{\frac{1}{2}-s}(\{|\tilde{a}_{i1}|^2\}) - \theta_{\frac{1}{2}-s}(K) \right| < \epsilon$ and $\left| \theta_{\frac{1}{2}+s}(\{|\tilde{a}_{i1}|^2\}) - \theta_{\frac{1}{2}+s}(K) \right| < \epsilon$ with probability $1 - \exp(-cm\epsilon^2)$. Thus, let $\zeta_s := \theta_{\frac{1}{2}-s}(K)$ and $\zeta^s := \theta_{\frac{1}{2}+s}(K)$, we have with probability $1 - \exp(-cm\epsilon^2)$

$$(\zeta_s - \epsilon)\|\mathbf{x}\|^2 < \theta_{1/2}(\{y_i\}) < (\zeta^s + \epsilon)\|\mathbf{x}\|^2, \quad (18)$$

where ζ^s and ζ_s can be arbitrarily close if s is small enough.

2. We next estimate the direction of \mathbf{x} , assuming $\|\mathbf{x}\| = 1$. The norm $\|\mathbf{x}\|$ can be estimated as in Algorithm 1. For simplicity of presentation, we assume $\|\mathbf{x}\| = 1$. Using (18), the matrix \mathbf{Y} in (7) can be bounded by $\mathbf{Y}_1 \prec \mathbf{Y} \prec \mathbf{Y}_2$ with high probability, where

$$\begin{aligned} \mathbf{Y}_1 &:= \frac{1}{m} \sum \mathbf{a}_i \mathbf{a}_i^T (\mathbf{a}_i^T \mathbf{x})^2 \mathbb{1}_{\{(\mathbf{a}_i^T \mathbf{x})^2 \leq \alpha_y^2 (\zeta_s - \epsilon)/0.455\}} \\ \mathbf{Y}_2 &:= \frac{1}{m} \sum \mathbf{a}_i \mathbf{a}_i^T (\mathbf{a}_i^T \mathbf{x})^2 \mathbb{1}_{\{(\mathbf{a}_i^T \mathbf{x})^2 \leq \alpha_y^2 (\zeta^s + \epsilon)/0.455\}} \\ &\quad + \frac{1}{m} \sum_{i \in \mathcal{S}} \mathbf{a}_i \mathbf{a}_i^T \alpha_y^2 (\zeta^s + \epsilon)/0.455. \end{aligned}$$

It can then be shown by concentration of random matrices with non-isotropic sub-Gaussian rows (Vershynin, 2012) that \mathbf{Y}_1 and \mathbf{Y}_2 concentrate on their means which can be made arbitrarily close with s sufficiently small. Together with the fact that both means of \mathbf{Y}_1 and \mathbf{Y}_2 have leading eigenvector \mathbf{x} , one can justify that the leading eigenvector of \mathbf{Y} can be made close enough to \mathbf{x} for sufficiently small constant s .

5.3. Geometric Convergence

We now show that within a small neighborhood of the ground truth, the truncated gradient (9) satisfies the *Regularity Condition (RC)* (Candès et al., 2015; Chen & Candès, 2015), which guarantees the geometric convergence of median-TWF once the initialization lands into this neighborhood.

Definition 2. The gradient $\nabla \ell_{tr}(\mathbf{z})$ is said to satisfy the *Regularity Condition* $\text{RC}(\mu, \lambda, \epsilon)$ if

$$-\left\langle \frac{1}{m} \nabla \ell_{tr}(\mathbf{z}), \mathbf{h} \right\rangle \geq \frac{\mu}{2} \left\| \frac{1}{m} \nabla \ell_{tr}(\mathbf{z}) \right\|^2 + \frac{\lambda}{2} \|\mathbf{h}\|^2 \quad (19)$$

for all \mathbf{z} and $\mathbf{h} = \mathbf{z} - \mathbf{x}$ obeying $\|\mathbf{h}\| \leq c\|\mathbf{z}\|$.

We will first show that $\nabla \ell_{tr}(\mathbf{z})$ satisfies the RC for the noise-free case, and then extend it to model (1) with sparse outliers, thus establishing the global convergence of median-TWF in both cases.

The central step to establish the RC is to show that the sample median used in the truncation rule concentrates at the level $\|\mathbf{z} - \mathbf{x}\| \|\mathbf{z}\|$ as stated in the following proposition.

Proposition 1. If $m > c_0 n \log n$, then with probability at least $1 - c_1 \exp(-c_2 m)$, there exist constants β and β' such that

$$\beta \|\mathbf{z}\| \|\mathbf{h}\| \leq \text{med}(\{|\mathbf{a}_i^T \mathbf{x}|^2 - |\mathbf{a}_i^T \mathbf{z}|^2\}_{i=1}^m) \leq \beta' \|\mathbf{z}\| \|\mathbf{h}\|,$$

holds for all $\mathbf{z}, \mathbf{h} := \mathbf{z} - \mathbf{x}$ satisfying $\|\mathbf{h}\| < 1/11\|\mathbf{z}\|$.

We note that a similar property for the sample mean has been shown in (Chen & Candès, 2015) as long as the number of measurements m is $O(n)$. In fact, the median is much more challenging to handle due to its non-linearity, which also causes slightly more measurements compared to the sample mean.

We next briefly sketch how we exploit the properties of the median developed in Section 5.1 to show Proposition 1. First fix \mathbf{x} and \mathbf{z} satisfying $\|\mathbf{x} - \mathbf{z}\| < 1/11\|\mathbf{z}\|$. It can be shown that

$$|(\mathbf{a}_i^T \mathbf{x})^2 - (\mathbf{a}_i^T \mathbf{z})^2| = c |u_i v_i| \cdot \|\mathbf{h}\| \|\mathbf{z}\|,$$

where u_i and v_i are correlated $\mathcal{N}(0, 1)$ Gaussian random variables and $1.89 < c < 2$. Hence, Lemma 4 and Lemma 1 imply that

$$\begin{aligned} (0.65 - \epsilon) \|\mathbf{z}\| \|\mathbf{h}\| &\leq \text{med}(\{ |(\mathbf{a}_i^T \mathbf{x})^2 - (\mathbf{a}_i^T \mathbf{z})^2| \}) \\ &\leq (0.91 + \epsilon) \|\mathbf{z}\| \|\mathbf{h}\| \end{aligned} \quad (20)$$

for given \mathbf{x} and \mathbf{z} with high probability. Then applying the net covering argument we prove that (20) holds for all \mathbf{z} and \mathbf{x} with $\|\mathbf{z} - \mathbf{x}\| \leq \frac{1}{11}\|\mathbf{z}\|$. In particular, Lemma 2 is applied to bound the distance between the medians of two points on and off the net.

For the model in (1) with outliers, we show that $\text{med}(\{|y_i - (\mathbf{a}_i \mathbf{z})^2|\})$ continues to have property as in Proposition 1 even with the presence of a small constant portion of outliers. This can be accomplished by first observing

$$\begin{aligned} \theta_{\frac{1}{2}-s}(\{ |(\mathbf{a}_i^T \mathbf{x})^2 - (\mathbf{a}_i^T \mathbf{z})^2 | \}) &\leq \theta_{\frac{1}{2}}(\{|y_i - (\mathbf{a}_i^T \mathbf{z})^2|\}) \\ &\leq \theta_{\frac{1}{2}+s}(\{ |(\mathbf{a}_i^T \mathbf{x})^2 - (\mathbf{a}_i^T \mathbf{z})^2 | \}). \end{aligned} \quad (21)$$

using Lemma 3, and then extending (20) to quantiles $\theta_{\frac{1}{2}-s}$ and $\theta_{\frac{1}{2}+s}$ respectively for a small constant s . Taking all these together yields bounds for both sides of (21) at the level of $\|\mathbf{z}\| \|\mathbf{h}\|$.

Remark 2. With ϵ small enough, $\beta = 0.6, \beta' = 1$ is a valid choice for Proposition 1. We will set the algorithm parameters based on these two values.

5.4. Stability with Additional Dense Bounded Noise

Now, consider the model in (2) with both sparse outliers and dense bounded noise. We omit the analysis of the initialization step as it is similar to Section 5.2. We split our analysis of the gradient loop into two regimes.

Regime 1: Assume $c_4\|z\| \geq \|h\| \geq c_3(\|w\|_\infty/\|z\|)$. Lemma 3 implies

$$\begin{aligned} \theta_{\frac{1}{2}-s}(\{|\tilde{y}_i - (\mathbf{a}_i^T z)^2|\}) &\leq \text{med}(\{|y_i - (\mathbf{a}_i^T z)^2|\}) \\ &\leq \theta_{\frac{1}{2}+s}(\{|\tilde{y}_i - (\mathbf{a}_i^T z)^2|\}), \end{aligned} \quad (22)$$

where $\tilde{y}_i := (\mathbf{a}_i^T x)^2 + w_i$ i.e., measurements that are corrupted by only bounded noise. Moreover, Lemma 2 and assumption of the regime implies

$$\begin{aligned} &\left| \theta_{\frac{1}{2}+s}(\{|\tilde{y}_i - (\mathbf{a}_i^T z)^2|\}) - \theta_{\frac{1}{2}+s}(\{|\mathbf{a}_i^T x)^2 - (\mathbf{a}_i^T z)^2|\}) \right| \leq \|w\|_\infty, \\ &\left| \theta_{\frac{1}{2}-s}(\{|\tilde{y}_i - (\mathbf{a}_i^T z)^2|\}) - \theta_{\frac{1}{2}-s}(\{|\mathbf{a}_i^T x)^2 - (\mathbf{a}_i^T z)^2|\}) \right| \leq \|w\|_\infty. \end{aligned} \quad (23)$$

Therefore, combining the above inequalities with Proposition 1, we have

$$\beta\|x - z\|\|z\| \leq \text{med}(\{|y_i - (\mathbf{a}_i^T z)^2|\}) \leq \beta'\|x - z\|\|z\|,$$

implying that RC holds in Regime 1, and the error decreases geometrically in each iteration.

Regime 2: Assume $\|h\| \leq c_3(\|w\|_\infty/\|z\|)$. Since each update $\frac{\mu}{m}\nabla\ell_{tr}(z)$ is at most the order of $\|w\|_\infty/\|z\|$, the estimation error cannot increase by more than $(\|w\|_\infty/\|z\|)$ with a constant factor. Thus, one has

$$\text{dist}\left(z + \frac{\mu}{m}\nabla\ell_{tr}(z), x\right) \leq c_5(\|w\|_\infty/\|x\|) \quad (24)$$

for some constant c_5 . As long as $\|w\|_\infty/\|x\|^2$ is sufficiently small, it is guaranteed that $c_5(\|w\|_\infty/\|x\|) \leq c_4\|x\|$. If the iteration jumps out of Regime 2, it falls into Regime 1 described above.

6. Numerical Experiments

In this section, we provide numerical experiments to demonstrate the effectiveness of median-TWF, which corroborates with our theoretical findings. We first show that, in the noise-free case, our median-TWF performs similarly as TWF (Chen & Candès, 2015) for exact recovery. We set the parameters of median-TWF as specified in Section 3, and those of TWF as suggested in (Chen & Candès, 2015). Let the signal length n take values from 1000 to 10000 by a step size of 1000, and the ratio of the sample complexity to the signal length, m/n , take values from 2 to 6 by

a step size of 0.1. For each pair of (m, n) , we generate a signal $x \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{n \times n})$, and the measurement vectors $\mathbf{a}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{n \times n})$ i.i.d. for $i = 1, \dots, m$. For both algorithms, a fixed number of iterations $T = 500$ are run, and the trial is declared successful if $z^{(T)}$, the output of the algorithm, satisfies $\text{dist}(z^{(T)}, x)/\|x\| \leq 10^{-8}$. Figure 1 shows the number of successful trials out of 20 trials for both algorithms, with respect to m/n and n . It can be seen that for both algorithms, as soon as m is above $4n$, exact recovery is achieved for both algorithms. Around the phase transition boundary, the performance of median-TWF is slightly worse than that of TWF, which is possibly due to the inefficiency of median compared to mean in the noise-free case (Huber, 2011).




Figure 1. Phase transition of median-TWF and TWF for noise-free data: the gray scale of each cell $(m/n, n)$ indicates the number of successful recovery out of 20 trials.

We next examine the performance of median-TWF in the presence of sparse outliers. We compare the performance of median-TWF with not only TWF but also an alternative which we call the *trimean-TWF*, based on replacing the sample mean in TWF by the *trimmed mean*. More specifically, trimean-TWF requires knowing the fraction s of outliers so that samples corresponding to sm largest gradient values are removed, and truncation is then based on the mean of remaining samples. We fix the signal length $n = 1000$ and the number of measurements $m = 8000$. We assume each measurement y_i is corrupted with probability $s \in [0, 0.4]$ independently, where the corruption value $\eta_i \sim \mathcal{U}(0, \|\eta\|_\infty)$ is randomly generated from a uniform distribution. Figure 2 shows the success rate of exact recovery over 100 trials as a function of s at different levels of outlier magnitudes $\|\eta\|_\infty/\|x\|^2 = 0.1, 1, 10, 100$, for the three algorithms median-TWF, trimean-TWF and TWF.

From Figure 2, it can be seen that median-TWF allows exact recovery as long as s is not too large for all levels of outlier magnitudes, without any knowledge of the outliers, which validates our theoretical analysis. Unsurprisingly, TWF fails quickly even with very small fraction of outliers. No successful instance is observed for TWF when $s \geq 0.02$ irrespective of the value of $\|\eta\|_\infty$. Trimean-TWF does not exhibit as sharp phase transition as median-TWF, and in general underperforms our median-TWF, except when both $\|\eta\|_\infty$ and s gets very large, see Figure 2(d). This is because in this range with large outliers, knowing the fraction s of outliers provides substantial advantage for




Figure 2. Success rate of exact recovery with outliers for median-TWF, trimean-TWF, and TWF at different levels of outlier magnitudes.




Figure 3. The relative error with respect to the iteration count for median-TWF and TWF with both dense noise and sparse outliers, and TWF with only dense noise. Performance of median-TWF with both dense noise and sparse outliers is comparable to that of TWF without outliers. (a) and (b): Uniform noise with different levels; (c) Poisson noise.

trimean-TWF to eliminate them, while the sample median can be deviated significantly from the true median for large s . Moreover, it is worth mentioning that exact recovery is more challenging for median-TWF when the magnitudes of most outliers are comparable to the measurements, as in Figure 2(c). In such a case, the outliers are more difficult to exclude as opposed to the case with very large outlier magnitudes as in Figure 2(d); and meanwhile, the outlier magnitudes in Figure 2(c) are large enough to affect the accuracy heavily in contrast to the cases in Figure 2(a) and 2(b) where outliers are less prominent.

We now examine the performance of median-TWF in the presence of both sparse outliers and dense bounded noise. The entries of the dense bounded noise w is generated independently from $\mathcal{U}(0, \|w\|_\infty)$, with $\|w\|_\infty/\|\mathbf{x}\|^2 = 0.001, 0.01$ respectively. The entries of the outlier is then generated as $\eta_i \sim \|w\| \cdot \text{Bernoulli}(0.1)$ independently. Figure 3(a) and Figure 3(b) depict the relative error $\text{dist}(z^{(t)}, \mathbf{x})/\|\mathbf{x}\|$ with respect to the iteration count t , for uniform noise at different levels. It can be seen that median-TWF under outlier corruption clearly outperforms TWF under the same situation, and acts as if the outliers do not exist by achieving almost the same accuracy as TWF under no outliers. Moreover, the solution accuracy of median-TWF has 10 times gain from Figure 3(a) to Figure 3(b) as $\|w\|_\infty$ shrinks by 1/10, which corroborates Theorem 2 nicely.

Finally, we consider when the measurements are corrupted both by Poisson noise and outliers, which models photon detection in optical imaging applications. We generate each measurement as $y_i \sim \text{Poisson}(\langle \mathbf{a}_i, \mathbf{x} \rangle^2)$, for

$i = 1, \dots, m$, which is then corrupted with probability $s = 0.1$ by outliers. The entries of the outlier are obtained by first generating $\eta_i \sim \|\mathbf{x}\|^2 \cdot \mathcal{U}(0, 1)$ independently, and then rounding it to the nearest integer. Figure 3(c) depicts the relative error $\text{dist}(z^{(t)}, \mathbf{x})/\|\mathbf{x}\|$ with respect to the iteration count t , where again median-TWF under both Poisson noise and sparse outlier noise has almost the same accuracy as TWF under only Poisson noise.

7. Conclusions

In this paper, we proposed median-TWF, and showed that it allows exact recovery even with a constant proportion of arbitrary outliers for robust phase retrieval. This is in contrast to recently developed WF and TWF, which likely to fail under outlier corruptions. We anticipate that sample median can be applied to designing provably robust non-convex algorithms for other inference problems under sparse arbitrary corruptions. The techniques we develop here to analyze performance guarantee for median-based algorithms will be useful in those contexts as well.

Acknowledgements

We would like to thank the anonymous reviewers for their helpful comments. We also thank Yining Wang and Yu-Xiang Wang at Carnegie Mellon University for helpful discussion on statistical properties of the median. The research of Huishuai Zhang and Yingbin Liang is supported by the grant AFOSR FA9550-16-1-0077. The work of Y. Chi is supported in part by the grants NSF CCF-1422966, AFOSR FA9550-15-1-0205 and ONR N00014-15-1-2387.

References

- Anandkumar, A., Jain, P., Shi, Y., and Niranjan, U.N. Tensor vs matrix methods: Robust tensor decomposition under block sparse perturbations. *arXiv preprint arXiv:1510.04747*, 2015.
- Arora, S., Ge, R., Ma, T., and Moitra, A. Simple, efficient, and neural algorithms for sparse coding. *arXiv preprint arXiv:1503.00778*, 2015.
- Balan, R., Casazza, P., and Edidin, D. On signal reconstruction without phase. *Applied and Computational Harmonic Analysis*, 20(3):345–356, 2006.
- Bhatia, K., Jain, P., and Kar, P. Robust regression via hard thresholding. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 721–729, 2015.
- Cai, T. T., Li, X., and Ma, Z. Optimal rates of convergence for noisy sparse phase retrieval via thresholded wirtinger flow. *arXiv preprint arXiv:1506.03382*, 2015.
- Candès, E. J. and Li, X. Solving quadratic equations via phaselift when there are about as many equations as unknowns. *Foundations of Computational Mathematics*, 14(5):1017–1026, 2014.
- Candès, E. J., Strohmer, T., and Voroninski, V. Phaselift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics*, 66(8):1241–1274, 2013.
- Candès, E. J., Li, X., and Soltanolkotabi, M. Phase retrieval via wirtinger flow: Theory and algorithms. *IEEE Transactions on Information Theory*, 61(4):1985–2007, 2015.
- Charikar, M., Chen, K., and Farach-Colton, M. Finding frequent items in data streams. In *Automata, languages and programming*, pp. 693–703. Springer, 2002.
- Chen, K. On k-median clustering in high dimensions. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp. 1177–1185. 2006.
- Chen, Y. and Candès, E. J. Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Advances in Neural Information Processing Systems (NIPS)*, pages 739–747. 2015.
- Chen, Y. and Wainwright, M. J. Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees. *arXiv preprint arXiv:1509.03025*, 2015.
- Chen, Y., Caramanis, C., and Mannor, S. Robust sparse regression under adversarial corruption. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pp. 774–782, 2013.
- Chen, Y., Yi, X., and Caramanis, C. A convex formulation for mixed regression with two components: Minimax optimal rates. In *Conf. on Learning Theory*, 2014.
- Chen, Y., Chi, Y., and Goldsmith, A.J. Exact and stable covariance estimation from quadratic sampling via convex programming. *IEEE Transactions on Information Theory*, 61(7):4034–4059, July 2015.
- Demanet, L. and Hand, P. Stable optimizationless recovery from phaseless linear measurements. *Journal of Fourier Analysis and Applications*, 20(1):199–221, 2014.
- Donahue, James D. Products and quotients of random variables and their applications. Technical report, DTIC Document, 1964.
- Drenth, J. *X-Ray Crystallography*. Wiley Online Library, 2007.
- Fienup, J. R. Phase retrieval algorithms: a comparison. *Applied optics*, 21(15):2758–2769, 1982.
- Hand, P. Phaselift is robust to a constant fraction of arbitrary errors. *arXiv preprint arXiv:1502.04241*, 2015.
- Huber, P. J. *Robust statistics*. Springer, 2011.
- Keshavan, R. H., Montanari, A., and Oh, S. Matrix completion from a few entries. *IEEE Transactions on Information Theory*, 56(6):2980–2998, June 2010.
- Lee, K., Li, Y., Junge, M., and Bresler, Y. Blind recovery of sparse signals from subsampled convolution. *arXiv preprint arXiv:1511.06149*, 2015.
- Li, X. and Voroninski, V. Sparse signal recovery from quadratic measurements via convex programming. *SIAM Journal on Mathematical Analysis*, 2013.
- Netrapalli, P., Niranjan, U.N., Sanghavi, S., Anandkumar, A., and Jain, P. Non-convex robust pca. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 1107–1115, 2014.
- Qu, C. and Xu, H. Subspace clustering with irrelevant features via robust dantzig selector. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 757–765. 2015.
- Soltanolkotabi, M. *Algorithms and theory for clustering and nonconvex quadratic programming*. PhD thesis, Stanford University, 2014.
- Sun, J., Qu, Q., and Wright, J. Complete dictionary recovery using nonconvex optimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 2351–2360, 2015.
- Tibshirani, R. J. Fast computation of the median by successive binning. *arXiv preprint arXiv:0806.3301*, 2008.

- Tu, S., Boczar, R., Soltanolkotabi, M., and Recht, B. Low-rank solutions of linear matrix equations via procrustes flow. *arXiv preprint arXiv:1507.03566*, 2015.
- Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. *Compressed Sensing, Theory and Applications*, pp. 210 – 268, 2012.
- Wagner, D. Resilient aggregation in sensor networks. In *Proceedings of the 2nd ACM workshop on Security of ad hoc and sensor networks*, pp. 78–87. ACM, 2004.
- Waldspurger, I., d’Aspremont, A., and Mallat, S. Phase recovery, maxcut and complex semidefinite programming. *Mathematical Programming*, 149(1-2):47–81, 2015.
- Weller, D.S., Pnueli, A., Divon, G., Radzyner, O., Eldar, Y.C., and Fessler, J.A. Undersampled phase retrieval with outliers. *IEEE Transactions on Computational Imaging*, 1(4):247–258, Dec 2015.
- White, C. D., Ward, R., and Sanghavi, S. The local convexity of solving quadratic equations. *arXiv preprint arXiv:1506.07868*, 2015.
- Zheng, Q. and Lafferty, J. A convergent gradient descent algorithm for rank minimization and semidefinite programming from random linear measurements. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 109–117, 2015.